Landscape Hardware Requirements

Kevin Retzke, Joe Boyd, Shreyas Bhat, Tanya Levshina

Introduction

Landscape Program has been introduced on SPPM meeting in October, 2015 and has been approved soon after. The goal of the program is to provide a comprehensive framework to allow the various teams in SCD, Fermilab projects and experiments to monitor services, jobs, data transfers, and storage usage on various facilities.
Currently this program collects data from GPGrid, HEP Cloud Facility, LPC CMS, Wilson Cluster and OSG sites. Accounting subproject of Landscape, GRACC, is scheduled to replace OSG accounting system, Gratia, in early March, 2017.

The program is getting more and more popular among our users and management. Landscape dashboards are used on daily basis by Fermilab users, sys administrators and support personnel. The program is a valuable tool for debugging various problems and analyzing trends. The plots have been used in many talks and presentations, including SPPM, CD All Hands meeting, Supercomputing and CHEP Conferences.

So far the Landscape team was able to utilize hardware by repurposing several nodes used for old Grid monitoring (fifemon) service and make informal agreements with Holland Computing Center (HCC) at Nebraska and the Grid Operation Center (GOC) at Indiana University.

We want to raise our concerns that current hardware is inadequate for Landscape goals. We also want to emphasis that we don’t have any formal agreements with HCC that is hosting ElasticSearch (ES) cluster or the GOC that allows us to use Rabbit MQ service for message exchange.

This document describes hardware requirement to make Landscape a state of the art program for Fermilab for the next 5 years.

Proposed Hardware Configuration

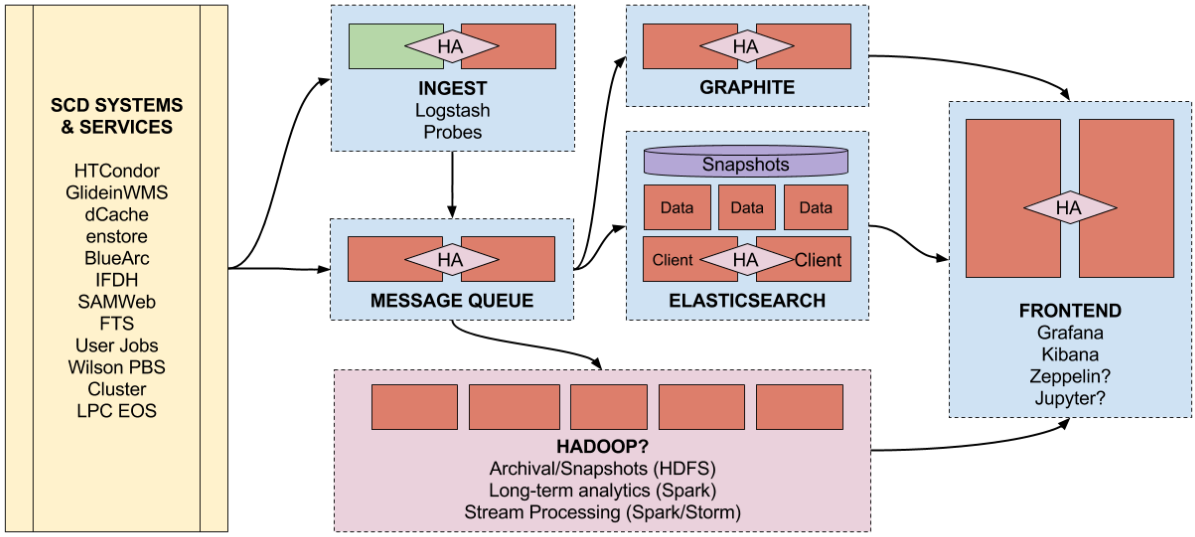

Detailed Hardware Specification

Services	# Nodes	# Core	Memory	Disk	Comments
ES Cluster	3	8	64 GB RAM	1 TB SSD data disk, separate system disk	Data and master nodes
	2	8	64 GB RAM	System disk	“Client” nodes that handle queries and web services (one prod and one dev/HA for prod)
				4TB shared filesystem	FS for snapshots for

					recovery for 2 years
Message Queue Service (RabbitMQ/Kafka)	2	8	64 GB RAM	1 TB SSD	Second node is needed for HA and dev.
Graphite	3	8	64 GB RAM	1 TB SSD	Two production nodes in full HA mode and one dev machine
Grafana	2	4	32 GB RAM		Could be VMs. Second node is needed for HA and dev uses)
Total	**10 servers + 2 VMs**	**88 cores**		**8 x 1 TB SSD drives 4TB shared file system**	

Hadoop Cluster, shown at the Hardware Configuration diagram in red, is a “reasonable” thing to have. It was included in the diagram because it is an essential component if we would like to follow the best practices analytics platform. We're not requesting it at this time.

ES Data Volume

Total volume of FIFE data currently is 300 GB which includes the replica. Data is replicated one master + one replica)

Currently collecting	**Retention**

Fifebatch HTCondor event logs	1 month
IFDH transfer logs	1 month
Fifebatch job and glideins classads	Indefinite

What else could we add?	**Estimate of space/cost (not included in "What we need" above)**
Job profiling	2TB for 1 month worth of data
System logs (from ECF)	Need to get from ECF
Glidein logs	100k x 2000 glideins x 2/day x 1 month =~ 100GB
HTCondor event logs from other clusters	2TB for GPGrid and CMS LPC cluster

HA and Dev Needs

Elasticsearch has built in redundancy for the data so the data nodes don't have to have duplicates.

Links to data for current Elasticsearch system

Disk usage is growing at 6.5 GB/day with GRACC data
https://gracc.opensciencegrid.org/dev/dashboard/db/grace-system-monitoring-disk-usage?from=now-30d&to=now
Disk I/O stats can be seen at
https://gracc.opensciencegrid.org/dev/dashboard/db/grace-system-monitoring-disk
General system metrics including network rate are at
https://gracc.opensciencegrid.org/dev/dashboard/db/grace-system-monitoring
Even more stats by node:
https://gracc.opensciencegrid.org/dev/dashboard/db/grace-system-monitoring-node-details
You can see a breakdown of space used by current indicies at:
https://gracc.opensciencegrid.org/dashboard/file/grace-indices.json

Message Queue (Rabbit MQ/Kafka)

Sending collected data through a message queue allows for flexibility in architecture. The data can be filtered and multiple consumers could subscribe to a stream (prod and dev for instance).

Graphite

2 nodes each: 8 cores, 64GB, 2TB SSD (prod with HA)
2 nodes each: 8 cores, 64GB, 2TB SSD (dev)
8 cores, 64GB, 2TB SSD/or platters for probe/snap/logstash machine (x2 for HA and dev, fifemon2 could be one of these it has warranty to 2020)

What else could we add	**Estimate of space/cost**
Basic info for offsite OSG clusters	10 sites x 5000 metrics
CMS Global pool stats	150,000 metrics

HA and Dev Need

Graphite has built in ability to replicate incoming data to a second server live.

Grafana

Could be a VM. 4 cores, 32 GB RAM, no special disk needs (x2 for dev and HA)

Current Grafana instances include:

1. FIFE: https://fifemon.fnal.gov/monitor/
2. LPC: https://landscape.fnal.gov/lpc/dashboard/db/lpc-summary
3. HEP Facility: https://fifemon.fnal.gov/hcf/dashboard/db/hep-cloud-demo
4. Wilson Cluster: https://fifemon-pp.fnal.gov/hpc/dashboard/db/wilson-cluster-usage
5. OSG GRACC: https://gracc.opensciencegrid.org/dashboard/db/osg-project-accounting

Active Archive Facility (AAF) data

Active Archive Facility (AAF) data that include enstore transfers and tape usage information and daily dCache storage snapshot are currently collected by dedicated Gratia Collector and stored in Gratia mysql database. The data is not sent to HCC ES cluster at Nebraska. Gratia will be decommissioned in March, 2017. If we want to preserve this data and continue gathering this information we will need additional XXXX.

OSG Funding issue

The current unofficial agreement with HCC, Nebraska allows us to store not only OSG related accounting records at HCC ES cluster but all logs related to FIFE activities. We have recently started to push LPC and Wilson Clusters data to this cluster. It is not clear for how long we will be able to do so. The data are not backed up. If we need to preserve all Fermilab accounting data in case of lack of OSG funding we will need additional 1TB SSD drive per ES node.

FY17 budget requests

We believe the FY17 should already have budget requests for the ES cluster and replacement machines for the current Fifemon service. The table of machines above include a request for additional machines to be used for a message queue service and one additional machine for a development Graphite installation. Separate VMs for the Grafana frontend are also additional. Depending on budgetary constraints, some existing machines could be reused and dev instances could be located on older hardware or virtualized.

Risk Assessment

During the last two years we have changed how monitoring is done at SCD. We have moved from home grown monitoring software, developed for each service, to the open source software stack that is used by hundreds well known companies, including EBay, Netflix, Cisco and others. The components of the stack are easily replaceable and we could swap any of them to more advanced open source software module. The ability to collect huge amount of logs and time series data in one place and do data mining and analysis of this information will allow us better understand resource usage trends, services performance and debug problems. The approach we have chosen provides the flexibility to handle collecting metrics from new and more diverse sources as needed to support the SCD mission.

We see the following issues with failure to acquire the specified hardware:

- Failure to have a local ES cluster will put the whole program at risk.
- Failure to replace current Graphite cluster will cause:
    - Long-term outage: fifemon1 is out of warrantly and fifemon2 will not be able to handle current load in the nearest future
    - Significant degradation of service, that have been already demonstrated by including data from the recent CMS Google test.
- Failure to provide hardware for HA and development environment will hinder the availability and reliability of the program.